arXiv:1609.07195v1 [stat.ME] 23 Sep 2016

Efficient Feature Selection With Large and High-dimensional Data

Néhémy Lim
Department of Statistics, University of Washington
and
Johannes Lederer
Departments of Statistics and Biostatistics, University of Washington

September 26, 2016

Abstract

Driven by the advances in technology, large and high-dimensional data have become the rule rather than the exception. Approaches that allow for feature selection with such data are thus highly sought after, in particular, since standard methods, like cross-validated Lasso, can be computationally intractable and, in any case, lack theoretical guarantees. In this paper, we propose a novel approach to feature selection in regression. Consisting of simple optimization steps and tests, it is computationally more efficient than existing methods and, therefore, suited even for very large data sets. Moreover, in contrast to standard methods, it is equipped with sharp statistical guarantees. We thus expect that our algorithm can help to leverage the increasing volume of data in Biology, Public Health, Astronomy, Economics, and other fields.

Keywords: high-dimensional regression, feature selection, tuning parameter, algorithms

1 Introduction

Feature selection is one of the most common statistical techniques to learn from data. In Systems Biology, just to name one field of application, feature selection is used to learn about the microbiological environment in the human gut from stool samples [1, 19]. The advent of Big Data provides large data sets that are expected to render feature selection possible at unprecedented resolutions. For this, however, major computational and statistical challenges yet need to be resolved.

A widely-used framework for feature selection is linear regression. Beyond being of interest by itself, it can serve as a basis for the analysis of network data [22] and data with instrumental variables [4]. Many contemporary data sets are "big" in the sense that both the number of parameters and the number of samples are large, and they are "high-dimensional" in the sense that the number of parameters can rival or even exceed the number of samples. On the other hand, one can often make the additional assumption that the true underlying model is "sparse" in the sense that there is a suitable model that is based on only a small number of parameters. For such data, the first choice is the families of penalized estimators, among them the Lasso [27], MCP [32], SCAD [13], Scaled Lasso [26] and Square-Root Lasso [5], and TREX [20]. However, these methods entail two main difficulties: First, even convex methods based on the Lasso or Scaled/Square-Root Lasso become infeasible in the presence of hundreds of thousands or even millions of samples and parameters, and much more so non-convex methods, such as MCP, SCAD, or TREX. Second, feature selection with penalized methods is hardly understood from a theoretical perspective. For example, the Lasso, MCP, and SCAD involve one or more tuning parameters, but the calibration of these tuning parameters with standard schemes such Cross-Validation, AIC, or BIC lacks any finite sample guarantees. On the other hand,

methods such as Square-Root/Scaled Lasso or TREX that aim at making tuning parameter selection superfluous, either require knowledge of model parameters that are inaccessible in practice, such as the noise distribution, or they lack a theoretical framework altogether.

In this paper, we introduce a novel approach to feature selection with large and high-dimensional data. Its two main properties are:

- It is equipped with sharp finite sample guarantees for arbitrary distributions of the noise and provides accurate feature selection in practice.
- It is computationally more efficient than standard approaches based on the Lasso, MCP, SCAD, or Scaled/Square-Root Lasso, allowing for the analysis of larger data sets.

After setting the framework and objective below, we introduce our approach in the form of an algorithm in Section 2. Two of its main ingredients are optimization steps of the Lasso objective function and AV_∞-tests [11] for calibration. The statistical guarantees are stated in Section 3, and the empirical performance of the method both in simulations and on real data is demonstrated in Section 4. We finally conclude with a discussion in Section 5. The proofs and further simulation results are deferred to the Appendix.

Framework and objective

We assume the linear regression model

$$Y = X\beta + \varepsilon\,,$$

where $Y \in \mathbb{R}^n$ is the outcome, $X \in \mathbb{R}^{n\times p}$ the design matrix, $\beta \in \mathbb{R}^p$ the regression vector, and $\varepsilon \in \mathbb{R}^n$ random noise. For simplicity, we assume that the design matrix is standardized, that is, for all $j \in \{1,\dots,p\}$, we assume $\sum_{i=1}^n X_{ij} = 0$ and $\sum_{i=1}^n X_{ij}^2 = n$. Importantly,

however, we allow for *arbitrary* noise distributions; in particular, we allow for correlated, heavy-tailed ε .

Our goal is feature selection, that is to estimate the support $S := \mathrm{supp}(\beta) := \{j : \beta_j \neq 0\}$. It is well known that the Lasso with optimal tuning parameter provides accurate feature selection under certain model assumptions [8, 17]. For vectors $\theta \in \mathbb{R}^p$ and tuning parameters $r > 0$, we thus consider the Lasso objective function [27]

$$f(\theta, r) := \|Y - X\theta\|_2^2 + r\|\theta\|_1 .$$

For given r, this provides the estimator $\hat{S}^r := \mathrm{supp}(\hat{\beta}^r)$, where

$$\hat{\beta}^r \in \operatorname*{argmin}_{\theta \in \mathbb{R}^p} f(\theta, r) . \tag{1}$$

In view of its - in a sense optimal - theoretical performance (see below), we consider the Lasso *with optimal tuning parameter* as our gold standard for feature selection. To be able to use this benchmark, we recall the model assumptions needed for successful feature selection with the Lasso. Recall that irrespective of the method, feature selection is feasible only if the correlations in the design matrix X are sufficiently small. Here, we impose two standard conditions on X. We first impose a sup-norm bound for the Lasso:

Condition 1.1. *Given any $\alpha > 0$, there exists a constant $c > 0$ and a (in general unknown) number r_α , the latter depending on α , such that for all $r \geq r_\alpha$, the two following conditions are met with probability at least $1 - \alpha$*

(i) $\hat{S}^r \subset S$;

(ii) $\|\beta - \hat{\beta}^r\|_\infty \leq cr/n$.

This condition is a condition on X in disguise: for example, Theorem 11.3 in the book [17] uses the primal-dual witness approach to derive a result of the above form under the irrepresentability condition on X. Further results of the above form with different assumptions

on X can be found in [9, 11, 21]. Imposing Condition 1.1 allows us to conveniently encompass all these results.

We also impose a version of the restricted eigenvalues condition itself:

Condition 1.2 (Restricted eigenvalues)**.** *For given $d \geq 0$, there is a constant $\gamma > 0$ such that*

$$\|X\delta\|_2^2 \geq n\gamma\|\delta\|_2^2$$

for all $\delta \in \mathbb{R}^p$ that satisfy $\|\delta_{S^c}\|_1 \leq 3\|\delta_S\|_1 + 2d/r$.

The restricted eigenvalues condition is standard in the statistics literature; see [8, 17] and references therein. The origin of the additive term $2d/r$ will become apparent later; in general, it is a small additive term since d is small and r will be of the order 1. We also refer to a related formulation in [23], which contains additive terms that originate in variations of the objective function. Finally, we note that the assumptions on X used to prove Condition 1.1, such as mutual coherence or irrepresentability, imply restricted eigenvalues [29]; however, we impose Condition 1.2 separately both for ease of presentation and to obtain sharper results.

Two important facts are that (i) the optimal tuning parameter is unknown in practice and (ii) Lasso computations can be very time-consuming or even infeasible for very large data sets. Hence, we seek a method for feature selection that has the same statistical guarantees as the (in practice unknown) set $\hat{S}^{r_\alpha}$ and that is computationally efficient.

2 Method

We now introduce our feature selection scheme. We call this method *FOS* (*F*ast and *O*ptimal *S*election) for convenient reference in the paper. A heuristic version of the method

is stated in Algorithm 1. The algorithm contains three main parts: First, for given tuning parameter, optimization steps are applied until the computational precision is sufficient. Second, a test determines the optimal stopping along the tuning parameter path. Third, an estimate of the support is computed by thresholding the current regression vector.

On a high level, our method is simply a roughly computed Lasso estimate with a subsequent thresholding of the elements. However, the difficult tasks are to answer the three following questions that correspond to the three parts described in the heuristic: What computational precision is required, and how to check if it is reached? How to test for an optimal stopping point along the tuning parameter path? What is the optimal value

Algorithm 1: Heuristics of FOS Scheme

Inputs : data (Y, X), sequence of M tuning parameters $r_1 = r_{\max} > \ldots > r_M > 0$

Output: estimated support

```
Initialize regression vector with some value;
Initialize tuning parameter with r_1;
repeat
    repeat
        Update regression vector using one pass of an optimization algorithm;
        Determine computational precision;
    until computational precision sufficient;
    Determine statistical precision;
    Take next tuning parameter in sequence;
until statistical precision sufficient;
Compute the output by thresholding the current value of the regression vector;
```

for thresholding? None of these questions is answered in the literature.

Algorithm 2 contains the precise version of the method. We now show how it solves the three tasks above. For this, we first disentangle two different kinds of quantities: We identify with hats $\hat{}$ quantities that refer to statistical estimators. Typically, these are merely theoretical quantities; for example, in finite time, we cannot exactly compute a Lasso solution $\hat{\beta}^r$. On the other hand, we identify with tildes $\tilde{}$ quantities that result from algorithms. These are the quantities that are of most practical relevance.

We use the concept of duality gap to ensure that the computational precision is sufficient. For given tuning parameter r, a dual formulation of the Lasso problem (1) reads [2]

$$\hat{\nu}^r := \operatorname*{argmax}_{\nu \in \mathbb{R}^n} \; D(\nu, r) \qquad \text{subject to } \|X^\top \nu\|_\infty \leq 1 \;, \tag{2}$$

where $D(\nu, r) := -r^2\|\nu + 2Y/r\|_2^2/4 + \|Y\|_2^2$. For any regression vector $\tilde{\beta}$ and any dual feasible variable $\tilde{\nu}$ of (2), the duality gap is $\Delta(\tilde{\beta}, \tilde{\nu}) := f(\tilde{\beta}, r) - D(\tilde{\nu}, r)$, and it is well known [7] that $\Delta(\tilde{\beta}, \tilde{\nu})$ is an upper bound of $f(\tilde{\beta}, r) - f(\hat{\beta}^r, r)$. This upper bound ensures that the required precision is reached. Importantly, we do not need to solve the dual problem of the Lasso, but instead, we only require a dual point, which can be found with an explicit expression. We refer to Appendix B for details.

The stopping point on the tuning parameter path is determined via AV_∞-tests [11].

Theorem 3.1 stated in the following section, more precisely, the bounds in sup-norm and the bound on the resulting tuning parameter, provide a proper value for the thresholding. Note that no similar results are known for Cross-Validation, BIC, AIC, or other standard calibration schemes, so that a theoretically justified thresholding procedure is not available for these methods.

Algorithm 2: FOS Scheme for Feature Selection in Linear Regression

Inputs : $Y \in \mathbb{R}^n$; $X \in \mathbb{R}^{n\times p}$; $r_1 = r_{\max} > r_2 > \ldots > r_M > 0$; $\gamma > 0$; $c > 0$

Output: $\tilde{S}$

Initialization : `statsCont:=true`; `statsIt:=1`; $\tilde{\beta}^{r_1} := 0$; $\widetilde{r} := r_M$;

while *statsCont==true AND statsIt* $< M$ **do**

  `statsIt:=statsIt+1`;

  `stopCrit:=false`;

  `betaOld:=`$\tilde{\beta}^{r_{\texttt{statsIt}-1}}$;

  while *stopCrit==false* **do**

    Compute a dual feasible point $\tilde{\nu}^{r_{\texttt{statsIt}}}$ of Problem (2);

    Compute the duality gap $\Delta(\tilde{\beta}^{r_{\texttt{statsIt}}}, \tilde{\nu}^{r_{\texttt{statsIt}}})$;

    if $\Delta(\tilde{\beta}^{r_{statsIt}}, \tilde{\nu}^{r_{statsIt}}) \leq \gamma c^2 r^2_{statsIt}/n$ **then**

      $\tilde{\beta}^{r_{\texttt{statsIt}}}$`:=betaOld`;

      `stopCrit:=true`;

    else

      $\tilde{\beta}^{r_{\texttt{statsIt}}}$`:=`$\mathcal{T}_{r_{\texttt{statsIt}}/L}\left(\texttt{betaOld} - 2X^\top(X\cdot\texttt{betaOld} - Y)/L\right)$;

      `betaOld:=`$\tilde{\beta}^{r_{\texttt{statsIt}}}$;

    end

  end

  `statsCont:=`$\prod_{k=1}^{\texttt{statsIt}} \mathbb{1}\left\{\|\tilde{\beta}^{r_{\texttt{statsIt}}} - \tilde{\beta}^{r_k}\|_\infty/(r_{\texttt{statsIt}} + r_k) - c/n \leq 0\right\}$;

end

if *statsCont==false* **then**

  $\widetilde{r} := r_{\texttt{statsIt}-1}$;

end

$\tilde{S}$:=$\{j \in \{1,\ldots,p\} \,:\, |\tilde{\beta}^{\widetilde{r}}_j| \geq 6c\widetilde{r}/n\}$

As initialization, we choose the all-zeros vector in $\mathbb{R}^p$, reflecting our assumption that many entries of the true regression vector are close to zero. As optimization algorithm, one could select proximal gradient descent, coordinate descent, or other techniques. We opted for the first one; the corresponding updates in Line 13 then read

$$\tilde{\beta} \mapsto \mathcal{T}_{\frac{r}{L}}(\tilde{\beta} - \frac{2}{L} X^\top (X\tilde{\beta} - Y)) \,,$$

where $\mathcal{T}$ is the elementwise soft-threshold operator defined by $\mathcal{T}_b(a)_j := \mathrm{sign}(a_j) \max(|a_j| - b, 0)$ for $j \in \{1, \ldots, p\}$, and where $L > 0$ is the step size determined by backtracking. In our data examples, we use the FISTA implementation [3] in `MATLAB`®.

Since we are limited to finitely many computations in practice, we consider finite sequences $r_1 = r_{\max} > r_2 > \cdots > r_M = r_{\min} > 0$. The concrete choice follows the ones used in standard implementations [15]: we use a logarithmically spaced grid of size $M = 100$, set $r_{\max} := 2\|X^\top Y\|_\infty$ to the smallest tuning parameter such that $\hat{\beta}^r = 0$, and define $r_{\min} := r_{\max}/u$ as a fraction of $r_{\max}$. Standard choices for u range from 100 to 10′000. On a very high level, $r_{\max}/r_\alpha \approx \|X^\top Y\|_\infty / \|X^\top \varepsilon\|_\infty \approx n\|\beta\|_1/\sqrt{n} \approx \sqrt{n}$. To ensure that $r_{\min} < r_\alpha$ on our data sets, we thus select $u := 1000$. Finally note that our theoretical results hold for any types of grids (also for continuous ranges of r), and because of the warm starts and the early stopping, the computational complexity of FOS depends only very mildly on M and $r_{\min}$.

Theory finally provides a precise guidance on the constants c and γ. Indeed, recall that c and γ are *not* tuning parameters, but instead dimensionless constants that specify the model assumptions. Therefore, for all practical purposes here, it is sufficient to hardcode the values $c = 0.75$ and $\gamma = 1$ that correspond to orthogonal design. (Note that in any case, feature selection needs the design to be nearly orthogonal [33].) However, for theoretical purposes, it is important to keep track of these constants. For example, estimators other

than the Lasso might verify Condition 1.1 with a different constant c. Our results can then be transferred directly using this value of c in the algorithm.

Computationally, FOS has two advantages: First, only a part of the tuning parameter path needs to be computed, more precisely, only the part with large and moderate tuning parameters. Second, only very rough computations are required; in particular, since a large tolerance can be accepted for large tuning parameters, only very small numbers of optimization cycles (in practice, often zero to five) are required per tuning parameter.

Statistically, our scheme has optimal accuracy for variable selection. Indeed, we prove that FOS provides the same statistical guarantees (up to a small factor) as if we computed the Lasso to convergence with optimal, but in practice unknown, tuning parameter.

In the following section, we establish the statistical optimality of our approach, see Theorem 3.1. After that, we demonstrate its empirical performance.

3 Statistical guarantees

Feature selection is a common statistical technique; yet, standard approaches based on the Lasso, SCAD, or MCP lack finite sample guarantees. The reason is that these approaches involve one or more tuning parameters, and corresponding calibration schemes such as Cross-Validation are not equipped with corresponding theory. Similarly, bounds for feature selection with methods such as the Square-Root Lasso, see [10], presume knowledge of inaccessible model parameters such as the noise distribution.

In strong contrast, our main result states that FOS is, up to small constant, equipped with the same statistical guarantees as the Lasso with optimal, in practice unknown tuning parameter. Indeed, we have the following.

Theorem 3.1. *For any $\alpha > 0$, it holds with probability at least $1 - \alpha$ that*

$$\|\beta - \tilde{\beta}^{\tilde{r}}\|_\infty \leq \frac{12cr_\alpha}{n}\,.$$

Moreover, it holds with probability at least $1-\alpha$ that $\tilde{r} \leq 2r_\alpha$ and, if $\min_{j\in S}|\beta_j| > 12cr_\alpha/n$, that

$$\tilde{S} \supset S\,.$$

This result proves two features of FOS: (1) False negative control: FOS recovers all sufficiently large coefficients. (2) Entry-wise control: the coefficient estimates associated with FOS are accurate; in particular, any coefficients corresponding to false positives are small. Note that the result holds for any $\alpha > 0$, while α does *not* need to be specified in FOS. Note also that, opposed to the Square-Root Lasso/Scaled Lasso approach, the noise distribution does *not* need to be known beforehand.

The optimality of the scheme can be seen when comparing the ℓ_∞-bound with the corresponding bound in Condition 1.1. Indeed, our scheme satisfies the same bound (up to the factor 12) as the Lasso with optimal tuning parameter.

We defer the proof of Theorem 3.1 to the Appendix. It consists of three parts: First, we derive an oracle inequality for vectors that are close to a Lasso solution. Second, we use Fenchel duality [7, 24] to show that our scheme is sufficiently accurate. Finally, we show that our application of the AV_∞-testing scheme [11] selects an optimal tuning parameter.

4 Empirical performance

We now demonstrate the computational efficiency and the empirical accuracy of FOS. To obtain a comprehensive overview, we consider a variety of settings, including synthetic data as well as biological and financial applications.

In Section 4.2, we show the scalability of our method by analyzing a financial data set with more than 150'000 parameters. The application to even larger data is currently limited by the memory restrictions in MATLAB®; a future C/Fortran implementation could remove this limitation. Equally important, however, is that the complexity often stems not from the size of a single data set but instead from the need for analyzing a large number of data sets. In Section 4.3, for example, we learn a biological network with a neighborhood selection scheme. Each of the corresponding regressions comprises only 500 samples and 1000 parameters. However, since 1000 such regressions are needed, the computational complexity can easily render standard methods infeasible. Therefore, the efficiency in regressions with moderate data is also a main concern; we analyze such settings in Section 4.1, where we invoke synthetic data with up to 10'000 samples and parameters.

We mainly compare FOS to Lasso with Cross-Validation (LassoCV), because the latter is currently the most popular method. However, we found the same conclusions when comparing to other Lasso-based approaches, such as Lasso combined with AIC or BIC. In particular, we rival or outmatch the feature selection performance of Lasso combined with BIC, a method specifically designed for this task. We refer to the Appendix for corresponding simulation results. On the other hand, non-convex approaches are presently not suited for our purposes. There has been recent progress regarding the computation of non-convex methods. For example, Wang et al. [30] show that path following algorithms can lead to more efficient computations of SCAD, MCP, and other non-convex methods, and Bien et al. [6] show how to solve a specific non-convex problem as a sequence of convex problems. However, despite such progress, non-convex methods remain computationally intractable for our data applications.

All computations are conducted with MATLAB® and are run on an E5-2640 Intel® Xeon® CPU (2.50GHz). FOS is implemented using the SPAMS package [2] coded in C++.

We compare with two LassoCV implementations: First, we implement LassoCV analogously to FOS using the `SPAMS` package and a 10-fold Cross-Validation with warm starts. This implementation, called $\text{LassoCV}_{\texttt{SPAMS}}$ in the following, is the most appropriate one for comparisons with the FOS implementation. However, much work has gone into efficient implementations of LassoCV. Therefore, we also use the well-known `glmnet` package [15] and call the corresponding implementation $\text{LassoCV}_{\texttt{glmnet}}$. However, these results must be treated with reservation, because `glmnet` cannot be calibrated to the same convergence criterion as our implementation. More precisely, the convergence criterion in `glmnet` needs to be specified in terms of maximum change in the objective, which does not coincide with the criterion in our algorithm and with the convergence of the estimator itself as needed in the theory. One could also argue that comparing FOS with $\text{LassoCV}_{\texttt{glmnet}}$ is not fair in any case, because `glmnet` exploits additional geometric properties of the Lasso (such as screening rules). These additional properties could also be used in our scheme, but their implementation is deferred to future work. However, we demonstrate that even in its current version, FOS can outperform both $\text{LassoCV}_{\texttt{SPAMS}}$ and $\text{LassoCV}_{\texttt{glmnet}}$.

4.1 Synthetic data

In this section, we use synthetic data to demonstrate the empirical performance of FOS. To this end, we generate data of two different sizes from linear regression models with $n = 500$ and $p = 1000$ and with $n = 5000$ and $p = 10'000$, respectively. More specifically, we sample each row of the design matrix $X \in \mathbb{R}^{n \times p}$ independently from a p-dimensional normal distribution with mean 0 and covariance matrix $(1 - \rho)\, \mathrm{I}_{p \times p} + \rho \mathbb{1}_{p \times p}$, where $\mathrm{I}_{p \times p}$ is the identity matrix, $\mathbb{1}_{p \times p}$ is the matrix of ones, and $\rho = 0.3$ is the correlation among the variables. We then normalize the matrix X so that its columns have Euclidean norm equal to $\sqrt{n}$. The entries of the noise $\varepsilon \in \mathbb{R}^n$ are generated according to a one-dimensional

standard normal distribution. The entries of β are first set to 0 except for 10 uniformly at random chosen entries that are each set to 1 or -1 with equal probability. The entire vector β is then rescaled such that the signal-to-noise ratio $\|X\beta\|_2^2/n$ is equal to 5.

We summarize the results in Table 1. The computational efficiency is measured in average timing in seconds; the statistical accuracy is measured in average Hamming distance, which is the sum of the number of false positives and the number of false negatives. We observe that FOS outperforms LassoCV$_{\texttt{SPAMS}}$ and LassoCV$_{\texttt{glmnet}}$ both in computational efficiency and in statistical accuracy. While we restrict the presentation to two settings, we found the same conclusion for a wide spectrum of parameters; we refer to the Appendix for additional results.

Table 1: Average run times (in seconds) and average Hamming distances for LassoCV$_{\texttt{SPAMS}}$, LassoCV$_{\texttt{glmnet}}$, and FOS. For the larger data set, LassoCV$_{\texttt{SPAMS}}$ timed out on our machine, which means that it took more than one hour on average.

	$n = 500$, $p = 1000$		$n = 5000$, $p = 10'000$	
Method	Timing	Hamming distance	Timing	Hamming distance
LassoCV$_{\texttt{SPAMS}}$	137.15 ± 9.33	56.00 ± 18.37	NA	NA
LassoCV$_{\texttt{glmnet}}$	2.08 ± 0.32	44.90 ± 16.07	111.47 ± 1.46	56.50 ± 23.93
FOS	0.10 ± 0.06	8.60 ± 3.10	4.81 ± 3.60	2.30 ± 4.16

An illustration of the two computational benefits of FOS is given in Figure 1. First, we observe that even with warm starts, LassoCV$_{\texttt{SPAMS}}$ requires a large number of iterations to converge. In contrast, FOS allows for early stopping, in particular, for large tuning parameters (recall that the required precision for FOS is proportional to the tuning parameter; instead, the required precision for other methods is unknown). Moreover, Cross-Validation,

BIC, AIC, and similar calibration schemes are based on the entire Lasso path, while only a part of the path is required for FOS.

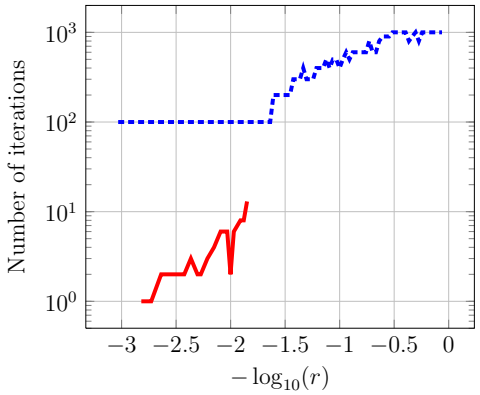

Figure 1: The red, solid line depicts the number of proximal gradient steps in FOS as a function of the tuning parameter r. The blue, dashed line depicts the corresponding number of proximal gradient steps in LassoCV$_{\texttt{SPAMS}}$. Shown are the numbers for one data set in the $n = 500$ and $p = 1000$ setting.

Another feature of FOS is the theoretically justified threshold. In contrast, there is no sound guidance for how to threshold LassoCV. This leads to many (tentatively small) false positives for LassoCV. Also, note that the results for LassoCV$_{\texttt{SPAMS}}$ and LassoCV$_{\texttt{glmnet}}$ differ, since their implementations are based on different algorithms and different stopping criteria.

4.2 Financial data

We now consider a large data set to demonstrate the scalability of FOS. The data [18] comprises $n = 16'087$ samples and $p = 150'348$ predictors and is publicly available.[1] The goal is to use financial reports of companies to predict the volatility of stock returns. The

[1] see `https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/regression.html`

feature representation of the financial reports is based on the calculation of TF-IDF (term frequency and inverse document frequency) of unigrams. The computational time (note that no ground truth is available for this data) of $\text{LassoCV}_{\texttt{glmnet}}$ is **153.02s** and of FOS **1.49s**. Instead, $\text{LassoCV}_{\texttt{SPAMS}}$ timed out on our machine.

4.3 Biological data

FOS can also be applied to structure learning problems by estimating the local neighborhood of each node via high-dimensional regressions. In this specific application [16], the goal is to understand the interaction network of $p = 1000$ genes in lung cancer patients from $n = 500$ expression profiles. We do neighborhood selection with the "or-rule" [22] based on FOS and LassoCV and compare the estimated graphs with the available gold standard [25]. The results are summarized in Figure 2. Note that here, the Hamming distance is the sum of the falsely included edges and the falsely omitted edges.

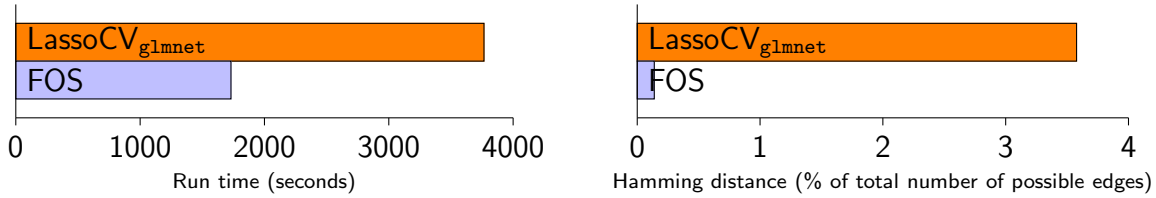

Figure 2: Run times (in seconds) and Hamming distances (in % of the total number of possible edges) for $\text{LassoCV}_{\texttt{glmnet}}$ and for FOS on the lung cancer data set. The implementation $\text{LassoCV}_{\texttt{SPAMS}}$ timed out on our machine.

5 Discussion

In view of the theoretical and empirical evidence provided above, FOS is a competitive approach to feature selection with large and high-dimensional data. The underlying theoretical results are precise guarantees for

- the computational accuracy needed along the tuning parameter path;
- the stopping criterion for the tuning parameter selection;
- the thresholding value.

For standard methods, no comparable guarantees are available.

Of interest for further research are screening rules for FOS. For the Lasso path, a variety of such rules have been developed and included in popular software such as `glmnet` [12, 14, 28, 31].

Another direction for further research is the application to other "base" estimators. In this study, we have focused on the Lasso as starting point. However, other methods, such as SCAD or MCP, satisfy similar ℓ_∞-bounds [23], and thus, could also fit our framework.

Finally, it would be of interest to study further the hypotheses needed for the theoretical results. Currently, the theory relies on strict assumptions on the design matrix. These assumptions are due to our focus on ℓ_∞-estimation and support recovery, which are possible only under strict conditions [33]. However, it would be interesting to consider extensions to other tasks, such as ℓ_2-estimation and prediction, that are possible under weaker assumptions on the design. Moreover, since the correlations are assumed small in our context, our theoretical and empirical results allow us to set $c = 0.75$ and $\gamma = 1$. At this point, however, it is unclear whether these are still appropriate choices if the correlations are large.

References

[1] AmGut. The American Gut Project. `http://humanfoodproject.com/americangut/`. Accessed: 2016-05-08.

[2] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Convex optimization with sparsity-inducing norms. In S. W. S. Sra, S. Nowozin, editor, *Optimization for Machine Learning*. MIT Press, 2011.

[3] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.*, 2(1):183–202, 2009.

[4] A. Belloni, V. Chernozhukov, and C. Hansen. Lasso methods for gaussian instrumental variables models. *arXiv:1012.1297*, 2011.

[5] A. Belloni, V. Chernozhukov, and L. Wang. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011.

[6] J. Bien, I. Gaynanova, J. Lederer, and C. Müller. Non-convex global minimization and false discovery rate control for the trex. *arXiv:1604.06815*, 2016.

[7] J. Borwein and A. Lewis. *Convex analysis and nonlinear optimization: theory and examples*. Springer-Verlag New York, 2010.

[8] P. Bühlmann and S. van de Geer. *Statistics for high-dimensional data: Methods, theory and applications*. Springer Ser. Statist. 2011.

[9] F. Bunea. Honest variable selection in linear and logistic regression models via ℓ_1 and $\ell_1 + \ell_2$ penalization. *Electron. J. Stat.*, 2:1153–1194, 2008.

[10] F. Bunea, J. Lederer, and Y. She. The Group Square-Root Lasso: Theoretical Properties and Fast Algorithms. *IEEE Trans. Inform. Theory*, 60(2):1313–1325, 2014.

[11] M. Chichignoud, J. Lederer, and M. Wainwright. Tuning Lasso for sup-norm optimality. *arXiv:1410.0247*, 2014.

[12] L. El Ghaoui, V. Viallon, and T. Rabbani. Safe feature elimination for the lasso and sparse supervised learning problems. *arXiv:1009.4219*, 2010.

[13] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.*, 96(456):1348–1360, 2001.

[14] O. Fercoq, A. Gramfort, and J. Salmon. Mind the duality gap: safer rules for the lasso. In D. Blei and F. Bach, editors, *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 333–342, 2015.

[15] J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22, 2010.

[16] I. Guyon, C. Aliferis, G. Cooper, A. Elisseeff, J.-P. Pellet, P. Spirtes, and A. Statnikov. Design and analysis of the causation and prediction challenge. In *WCCI Causation and Prediction Challenge*, pages 1–33, 2008.

[17] T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman and Hall, 2015.

[18] S. Kogan, D. Levin, B. Routledge, J. Sagi, and N. Smith. Predicting risk from financial reports with regression. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 272–280, 2009.

[19] Z. Kurtz, C. Müller, E. Miraldi, D. Littman, M. Blaser, and R. Bonneau. Sparse and compositionally robust inference of microbial ecological networks. *PLoS Comput. Biol.*, 11(5):e1004226, 2015.

[20] J. Lederer and C. Müller. Don't fall for tuning parameters: Tuning-free variable selection in high dimensions with the TREX. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 2729–2735, 2015.

[21] K. Lounici. Sup-norm convergence rate and sign concentration property of Lasso and Dantzig estimators. *Electron. J. Stat.*, 2:90–102, 2008.

[22] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the lasso. *Ann. Statist.*, 34(3):1436–1462, 2006.

[23] S. Negahban, P. Ravikumar, M. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statist. Sci.*, 27(4):538–557, 2012.

[24] R. Rifkin and R. Lippert. Value regularization and fenchel duality. *J. Mach. Learn. Res.*, 8:441–479, 2007.

[25] A. Statnikov, S. Ma, M. Henaff, N. Lytkin, E. Efstathiadis, E. Peskin, and C. Aliferis. Ultra-scalable and efficient methods for hybrid observational and experimental local causal pathway discovery. *J. Mach. Learn. Res.*, 16:3219–3267, 2015.

[26] T. Sun and C. Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012.

[27] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288, 1996.

[28] R. Tibshirani, J. Bien, J. Friedman, T. Hastie, N. Simon, J. Taylor, and R. Tibshirani. Strong rules for discarding predictors in lasso-type problems. *J. Roy. Statist. Soc. Ser. B*, 74(2):245–266, 2012.

[29] S. van de Geer and P. Bühlmann. On the conditions used to prove oracle results for the Lasso. *Electron. J. Stat.*, 3:1360–1392, 2009.

[30] Z. Wang, H. Liu, and T. Zhang. Optimal computational and statistical rates of convergence for sparse nonconvex learning problems. *Ann. Statist.*, 42(6):2164, 2014.

[31] Z. Xiang, Y. Wang, and P. Ramadge. Screening tests for lasso problems. *arXiv:1405.4897*, 2014.

[32] C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.*, pages 894–942, 2010.

[33] P. Zhao and B. Yu. On model selection consistency of Lasso. *J. Mach. Learn. Res.*, 7:2541–2563, 2006.

[34] H. Zou, T. Hastie, and R. Tibshirani. On the "degrees of freedom" of the lasso. *Ann. Statist.*, 35(5):2173–2192, 2007.

APPENDIX

In the first three parts of the Appendix, we introduce the building blocks for proving the claim stated in the main part of the paper, Theorem 3.1. Specifically, in Appendix A, we prove a geometric property (Lemma A.1) and an oracle inequality (Theorem A.1) for approximate Lasso solutions. Next, we show in Appendix B how to verify if the required computational precision is reached (Theorem 3.1). Using theses results, we then prove our main result (Theorem 3.1) in Appendix C. The dependencies among the lemmas and theorems are depicted in Figure 3. In the final part of the Appendix, Appendix D, we present further simulation results.

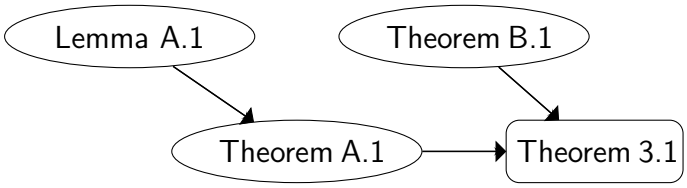

Figure 3: Dependencies among the lemmas and theorems. For example, the arrow between Lemma A.1 and Theorem A.1 indicates that Theorem A.1 relies on Lemma A.1.

A Underlying statistical results

In this section, we derive two properties of all vectors close to a Lasso solution: In Lemma A.1, we show that their error belongs to a cone; in Theorem A.1, we show that they satisfy an oracle inequality in ℓ_∞-norm.

Lemma A.1. *Let $d \geq 0$ be a constant, $r \geq 2r_\alpha$ a tuning parameter, and $\tilde{\beta} \in \mathbb{R}^p$ any vector that satisfies $f(\tilde{\beta}, r) \leq f(\hat{\beta}^r, r) + d$. Then $\delta := \hat{\beta}^r - \tilde{\beta}$ belongs to the cone*

$$\mathcal{C}(S) := \left\{ \nu \in \mathbb{R}^p \,:\, \|\nu_{S^c}\|_1 \leq 3\|\nu_S\|_1 + \frac{2d}{r} \right\}.$$

Proof of Lemma A.1. Since $f(\tilde{\beta}, r) \leq f(\hat{\beta}^r, r) + d \leq f(\hat{\beta}^{r/2}, r) + d$, we find the basic inequality

$$\|Y - X\tilde{\beta}\|_2^2 + r\|\tilde{\beta}\|_1 \leq \|Y - X\hat{\beta}^{r/2}\|_2^2 + r\|\hat{\beta}^{r/2}\|_1 + d\,.$$

We can now rewrite $\|Y - X\tilde{\beta}\|_2^2$ as follows:

$$\begin{aligned}\|Y - X\tilde{\beta}\|_2^2 &= \|Y - X\hat{\beta}^{r/2} + X\hat{\beta}^{r/2} - X\tilde{\beta}\|_2^2 \\ &= \|Y - X\hat{\beta}^{r/2}\|_2^2 + 2\,\langle Y - X\hat{\beta}^{r/2},\, X\hat{\beta}^{r/2} - X\tilde{\beta}\rangle + \|X\hat{\beta}^{r/2} - X\tilde{\beta}\|_2^2 \\ &= \|Y - X\hat{\beta}^{r/2}\|_2^2 + 2\,\langle X^\top(Y - X\hat{\beta}^{r/2}),\, \hat{\beta}^{r/2} - \tilde{\beta}\rangle + \|X\hat{\beta}^{r/2} - X\tilde{\beta}\|_2^2\,.\end{aligned}$$

Combining the two displays yields

$$2\,\langle X^\top(Y - X\hat{\beta}^{r/2}),\, \hat{\beta}^{r/2} - \tilde{\beta}\rangle + \|X\hat{\beta}^{r/2} - X\tilde{\beta}\|_2^2 + r\|\tilde{\beta}\|_1 \leq r\|\hat{\beta}^{r/2}\|_1 + d\,,$$

and, by rearranging and using that $\|X\hat{\beta}^{r/2} - X\tilde{\beta}\|_2^2 \geq 0$,

$$r\|\tilde{\beta}\|_1 \leq \langle 2X^\top(Y - X\hat{\beta}^{r/2}),\, \tilde{\beta} - \hat{\beta}^{r/2}\rangle + r\|\hat{\beta}^{r/2}\|_1 + d\,.$$

Invoking Hölder's inequality and the KKT conditions for $\hat{\beta}^{r/2}$ provides us with

$$\begin{aligned}2\,\langle X^\top(Y - X\hat{\beta}^{r/2}),\, \tilde{\beta} - \hat{\beta}^{r/2}\rangle &\leq \|2X^\top(Y - X\hat{\beta}^{r/2})\|_\infty \|\tilde{\beta} - \hat{\beta}^{r/2}\|_1 \\ &\leq \frac{r}{2}\|\tilde{\beta} - \hat{\beta}^{r/2}\|_1\,.\end{aligned}$$

Combining the two displays yields

$$2\|\tilde{\beta}\|_1 \leq \|\tilde{\beta} - \hat{\beta}^{r/2}\|_1 + 2\|\hat{\beta}^{r/2}\|_1 + \frac{2d}{r}\,.$$

Hence,

$$\begin{aligned}2\|\tilde{\beta}_{\hat{S}}\|_1 + 2\|\tilde{\beta}_{\hat{S}^c}\|_1 &\leq \|\tilde{\beta}_{\hat{S}} - \hat{\beta}^{r/2}_{\hat{S}}\|_1 + \|\tilde{\beta}_{\hat{S}^c} - \hat{\beta}^{r/2}_{\hat{S}^c}\|_1 + 2\|\hat{\beta}^{r/2}_{\hat{S}}\|_1 + 2\|\hat{\beta}^{r/2}_{\hat{S}^c}\|_1 + \frac{2d}{r} \\ &= \|\tilde{\beta}_{\hat{S}} - \hat{\beta}^{r/2}_{\hat{S}}\|_1 + \|\tilde{\beta}_{\hat{S}^c}\|_1 + 2\|\hat{\beta}^{r/2}_{\hat{S}}\|_1 + \frac{2d}{r}\,,\end{aligned}$$

where $\hat{S} := \mathrm{supp}(\hat{\beta}^{r/2})$. This is equivalent to

$$\|\tilde{\beta}_{\hat{S}^c}\|_1 \leq \|\tilde{\beta}_{\hat{S}} - \hat{\beta}^{r/2}_{\hat{S}}\|_1 + 2\|\hat{\beta}^{r/2}_{\hat{S}}\|_1 - 2\|\tilde{\beta}_{\hat{S}}\|_1 + \frac{2d}{r}\,,$$

so that with the reverse triangle inequality

$$\|\tilde{\beta}_{\hat{S}^c}\|_1 \leq 3\|\tilde{\beta}_{\hat{S}} - \hat{\beta}^{r/2}_{\hat{S}}\|_1 + \frac{2d}{r}\,.$$

Finally, setting $\delta := \hat{\beta}^r - \tilde{\beta}$, we get

$$\|\delta_{\hat{S}^c}\|_1 \leq 3\|\delta_{\hat{S}}\|_1 + \frac{2d}{r}\,.$$

Now, since $r/2 \geq r_\alpha$, Condition 1.1 (i) entails $\|\delta_{\hat{S}}\|_1 \leq \|\delta_S\|_1$ and $\|\delta_{S^c}\|_1 \leq \|\delta_{\hat{S}^c}\|_1$. Combining these two inequalities with the above display yields

$$\|\delta_{S^c}\|_1 \leq \|\delta_{\hat{S}^c}\|_1 \leq 3\|\delta_{\hat{S}}\|_1 + \frac{2d}{r} \leq 3\|\delta_S\|_1 + \frac{2d}{r}$$

as desired. □

Theorem A.1. *Suppose that $\alpha > 0$ and $r \geq 2r_\alpha$. Then, for any $d \geq 0$ and any $\tilde{\beta} \in \mathbb{R}^p$ with $f(\tilde{\beta}, r) \leq f(\hat{\beta}^r, r) + d$, it holds with probability at least $1 - \alpha$ that*

$$\|\beta - \tilde{\beta}\|_\infty \leq \frac{cr}{n} + \sqrt{\frac{d}{n\gamma}}\,.$$

Proof of Theorem A.1. In view of Condition 1.1 (ii), it is sufficient to show that

$$\|\hat{\beta}^r - \tilde{\beta}\|_\infty \leq \sqrt{\frac{d}{n\gamma}}\,.$$

Indeed, the desired claim follows directly from this via the triangle inequality for norms.

Let us prove the above inequality. Since $f(\hat{\beta}^r, r) \leq f(\tilde{\beta}, r) \leq f(\hat{\beta}^r, r) + d$, we find the basic equality

$$\|Y - X\tilde{\beta}\|_2^2 + r\|\tilde{\beta}\|_1 = \|Y - X\hat{\beta}^r\|_2^2 + r\|\hat{\beta}^r\|_1 + d'\,,$$

for some $d' \in [0, d]$. This equation is equivalent to

$$\|Y - X\hat{\beta}^r + X\hat{\beta}^r - X\tilde{\beta}\|_2^2 + r\|\tilde{\beta}\|_1 = \|Y - X\hat{\beta}^r\|_2^2 + r\|\hat{\beta}^r\|_1 + d' \,,$$

and hence to

$$\|Y - X\hat{\beta}^r\|_2^2 + 2\,\langle Y - X\hat{\beta}^r,\, X\hat{\beta}^r - X\tilde{\beta}\rangle + \|X\hat{\beta}^r - X\tilde{\beta}\|_2^2 + r\|\tilde{\beta}\|_1 = \|Y - X\hat{\beta}^r\|_2^2 + r\|\hat{\beta}^r\|_1 + d' \,,$$

and finally

$$\|Y - X\hat{\beta}^r\|_2^2 + \langle 2X^\top(Y - X\hat{\beta}^r),\, \hat{\beta}^r - \tilde{\beta}\rangle + \|X\hat{\beta}^r - X\tilde{\beta}\|_2^2 + r\|\tilde{\beta}\|_1 = \|Y - X\hat{\beta}^r\|_2^2 + r\|\hat{\beta}^r\|_1 + d' \,.$$

Rearranging yields

$$\|X\hat{\beta}^r - X\tilde{\beta}\|_2^2 = \langle 2X^\top(Y - X\hat{\beta}^r),\, \tilde{\beta} - \hat{\beta}^r\rangle - r\|\tilde{\beta}\|_1 + r\|\hat{\beta}^r\|_1 + d' \,.$$

We now recall that the KKT conditions for the objective function f read

$$-2X^\top(Y - X\hat{\beta}^r) + r\hat{\kappa} = 0$$

for any vector $\hat{\kappa}$ that satisfies $\|\hat{\kappa}\|_\infty \leq 1$ and $\hat{\kappa}^\top\hat{\beta}^r = \|\hat{\beta}^r\|_1$. Plugging this into the above display yields

$$\begin{aligned}\|X\hat{\beta}^r - X\tilde{\beta}\|_2^2 &= \langle r\hat{\kappa},\, \tilde{\beta} - \hat{\beta}^r\rangle - r\|\tilde{\beta}\|_1 + r\|\hat{\beta}^r\|_1 + d' \\ &= r(\hat{\kappa}^\top\tilde{\beta} - \|\tilde{\beta}\|_1) + d' \,,\end{aligned}$$

and by Hölder's inequality, $\hat{\kappa}^\top\tilde{\beta} \leq \|\hat{\kappa}\|_\infty\|\tilde{\beta}\|_1 \leq \|\tilde{\beta}\|_1$. Therefore,

$$\|X\hat{\beta}^r - X\tilde{\beta}\|_2^2 \leq d' \leq d \,.$$

We finally get, setting $\delta := \hat{\beta}^r - \tilde{\beta}$,

$$\|X\delta\|_2^2 \leq d \,.$$

Moreover, Lemma A.1 guarantees that $\delta \in \mathcal{C}(S)$ and thus allows us to apply Condition 2.1. So we also find

$$\|X\delta\|_2^2 \geq n\gamma\|\delta\|_2^2 .$$

Combining these two inequalities gives us

$$\|\delta\|_2^2 \leq \frac{d}{n\gamma},$$

which implies

$$\|\delta\|_\infty \leq \sqrt{\frac{d}{n\gamma}}$$

as desired. □

B Underlying optimization results

Our approach estimates the computational accuracies of iterative optimization steps. For this, we invoke Fenchel duality [7], which is used for many learning tasks in which the objective function can be split into a convex, differentiable fitting term and a convex, possibly non-smooth regularization term [24]. The Fenchel conjugate, which is central to Fenchel duality, is defined by

$$g^*(\eta) := \sup_{\omega \in \mathbb{R}^n} \left\{\eta^\top \omega - g(\omega)\right\}$$

for any function $g : \mathbb{R}^n \to [-\infty, \infty]$.

Theorem B.1. *Let $\tilde{\beta} \in \mathbb{R}^p$ be an estimate of $\hat{\beta}^r$. (i) A feasible point of the dual problem of the Lasso is*

$$\tilde{\nu} = \frac{2s}{r}(X\tilde{\beta} - Y)$$

where

$$s = \min\left\{\max\left\{\frac{-r}{\|2X^\top(X\tilde\beta - Y)\|_\infty}, \frac{-Y^\top(X\tilde\beta - Y)}{\|Y - X\tilde\beta\|_2^2}\right\}, \frac{r}{\|2X^\top(X\tilde\beta - Y)\|_\infty}\right\}$$

Moreover (ii) a duality gap is

$$\Delta(\tilde\beta,\, \tilde\nu) = \|Y - X\tilde\beta\|_2^2 + r\|\tilde\beta\|_1 + \frac{r^2}{4}\|\tilde\nu + \frac{2Y}{r}\|_2^2 - \|Y\|_2^2\,.$$

Proof of Theorem B.1. The Lasso problem is equivalent to the following primal optimization problem

$$\hat\beta^r \in \operatorname*{argmin}_{\theta\in\mathbb{R}^p} \quad \{g(\omega) + r\|\theta\|_1\}$$
$$\text{subject to} \quad \omega = X\theta\,,$$

where $g(\omega) := \|Y - \omega\|_2^2$ with Fenchel conjugate $g^*(\eta) = \|\eta + 2Y\|_2^2/4 - \|Y\|_2^2$. The corresponding Fenchel dual problem then reads [2]

$$\hat\nu^r := \operatorname*{argmax}_{\nu\in\mathbb{R}^n} \quad -g^*(r\nu)$$
$$\text{subject to} \quad \|X^\top\nu\|_\infty \le 1\,.$$

where $g^*(r\nu) = r^2\|\nu + 2Y/r\|_2^2/4 - \|Y\|_2^2$. A primal solution of the Lasso and the dual solution are linked by the relation

$$2(X\hat\beta^r - Y) = r\hat\nu^r\,.$$

Given a current primal estimate $\tilde\beta$, a dual feasible variable $\tilde\nu$ can be chosen as the closest (in ℓ_2-norm) point to $-2Y/r$ proportional to $2(X\tilde\beta - Y)/r$. This yields $\tilde\nu = 2s(X\tilde\beta - Y)/r$, where s is given by [12]

$$s = \min\left\{\max\left\{\frac{-r}{\|2X^\top(X\tilde\beta - Y)\|_\infty}, \frac{-Y^\top(X\tilde\beta - Y)}{\|Y - X\tilde\beta\|_2^2}\right\}, \frac{r}{\|2X^\top(X\tilde\beta - Y)\|_\infty}\right\}\,.$$

This concludes the first part of the proof.

Hence, a duality gap for the Lasso problem is

$$\Delta(\tilde{\beta},\, \tilde{\nu}) = g(X\tilde{\beta}) + r\|\tilde{\beta}\|_1 + g^*(r\tilde{\nu}) = \|Y - X\tilde{\beta}\|_2^2 + r\|\tilde{\beta}\|_1 + \frac{r^2}{4}\|\tilde{\nu} + \frac{2Y}{r}\|_2^2 - \|Y\|_2^2\,.$$

This concludes the second part of the proof. □

C Proof of Theorem 3.1

Having derived the oracle inequality stated in Theorem A.1 and the computational accuracy stated in Theorem B.1, we are now ready to prove our main result.

Proof of Theorem 3.1. Let $r \geq 2r_\alpha$ and $(\tilde{\beta}, \tilde{\nu})$ any pair of Lasso primal-dual variables that satisfies $\Delta(\tilde{\beta},\, \tilde{\nu}) \leq \gamma c^2 r^2/n$. According to our results above, it holds with probability at least $1-\alpha$ that

$$\|\beta - \tilde{\beta}\|_\infty \leq \frac{2cr}{n}\,.$$

Since our scheme invokes AV_∞-tests introduced in [11], we can now prove Theorem 3.1 along the same lines as [11, Theorem 1]. □

D Additional simulation results

We provide additional simulation results to complement the empirical studies in the main part of the paper. We consider the $n = 500$ and $p = 1000$ setting with the specifications as described in the main part of the paper - except for the stated differences.

In Table 2, we add the results for Lasso with AIC and BIC. In Table 3, we consider different support sizes. In Table 4, we vary the correlation.

The BIC and AIC criteria are defined as [34]

$$\mathrm{BIC}(r) = \frac{\|Y - X\tilde{\beta}^r\|_2^2}{n\sigma_\varepsilon^2} + \frac{\log(n)|\tilde{S}^r|}{n}$$
$$\mathrm{AIC}(r) = \frac{\|Y - X\tilde{\beta}^r\|_2^2}{n\sigma_\varepsilon^2} + \frac{2|\tilde{S}^r|}{n},$$

where σ_ε^2 is the variance of the noise ε. In our experiments, $\sigma_\varepsilon^2 = 1$. To establish valid comparisons, we again implement these schemes with the `SPAMS` package.

Table 2: Average run times (in seconds) and average Hamming distances for $\mathrm{LassoCV}_{\texttt{SPAMS}}$, $\mathrm{LassoCV}_{\texttt{glmnet}}$, $\mathrm{LassoBIC}_{\texttt{SPAMS}}$, $\mathrm{LassoAIC}_{\texttt{SPAMS}}$, and FOS.

Method	Timing	Hamming distance
$\mathrm{LassoCV}_{\texttt{SPAMS}}$	137.15 ± 9.33	56.00 ± 18.37
$\mathrm{LassoCV}_{\texttt{glmnet}}$	2.08 ± 0.32	44.90 ± 16.07
$\mathrm{LassoBIC}_{\texttt{SPAMS}}$	20.12 ± 5.98	10.40 ± 7.15
$\mathrm{LassoAIC}_{\texttt{SPAMS}}$	20.12 ± 5.98	45.30 ± 21.40
FOS	0.10 ± 0.06	8.60 ± 3.10

Table 3: Average run times (in seconds) and average Hamming distances for LassoCV$_{\texttt{SPAMS}}$, LassoCV$_{\texttt{glmnet}}$, and FOS for two different sizes of the true support: $|S| \in \{5, 30\}$.

	$\lvert S\rvert = 5$		$\lvert S\rvert = 30$	
Method	Timing	Hamming distance	Timing	Hamming distance
LassoCV$_{\texttt{SPAMS}}$	283.19 ± 75.41	43.50 ± 19.53	224.92 ± 41.77	117.50 ± 22.49
LassoCV$_{\texttt{glmnet}}$	2.59 ± 0.40	29.90 ± 12.66	2.52 ± 0.25	98.00 ± 35.50
FOS	0.76 ± 2.13	4.40 ± 1.58	2.23 ± 2.01	20.60 ± 11.98

Table 4: Average run times (in seconds) and average Hamming distances for LassoCV$_{\texttt{SPAMS}}$, LassoCV$_{\texttt{glmnet}}$, and FOS for two different strengths of the pairwise correlations: $\rho \in \{0, 0.4\}$.

	$\rho = 0$		$\rho = 0.4$	
Method	Timing	Hamming distance	Timing	Hamming distance
LassoCV$_{\texttt{SPAMS}}$	168.03 ± 25.28	56.50 ± 22.45	177.79 ± 48.11	58.30 ± 20.67
LassoCV$_{\texttt{glmnet}}$	2.46 ± 0.24	40.40 ± 13.17	2.36 ± 0.21	44.50 ± 18.12
FOS	1.78 ± 0.73	1.10 ± 1.29	0.12 ± 0.06	9.70 ± 0.95